arXiv:1703.09554v1 [cs.CV] 28 Mar 2017

Lucid Data Dreaming for Object Tracking

Anna Khoreva[1] Rodrigo Benenson[2] Eddy Ilg[3] Thomas Brox[3] Bernt Schiele[1]

[1]Max Planck Institute for Informatics, Germany
{khoreva, schiele}@mpi-inf.mpg.de

[2]Google
benenson@google.com

[3]University of Freiburg, Germany
{ilg, brox}@cs.uni-freiburg.de

Abstract

Convolutional networks reach top quality in pixel-level object tracking but require a large amount of training data ($1k\sim10k$) to deliver such results. We propose a new training strategy which achieves state-of-the-art results across three evaluation datasets while using $20\times\sim100\times$ less annotated data than competing methods.

Instead of using large training sets hoping to generalize across domains, we generate in-domain training data using the provided annotation on the first frame of each video to synthesize ("lucid dream"[1]) plausible future video frames. In-domain per-video training data allows us to train high quality appearance- and motion-based models, as well as tune the post-processing stage. This approach allows to reach competitive results even when training from only a single annotated frame, without ImageNet pre-training. Our results indicate that using a larger training set is not automatically better, and that for the tracking task a smaller training set that is closer to the target domain is more effective. This changes the mindset regarding how many training samples and general "objectness" knowledge are required for the object tracking task.

1. Introduction

In the last years the field of object tracking in videos has transitioned from bounding box [27, 29, 28] to pixel-level tracking [30, 44, 41, 57]. Given a first frame labelled with the foreground object mask, one aims to find the corresponding object pixels in future frames. Tracking objects at the pixel level enables a finer understanding of videos and is helpful for tasks such as video editing, rotoscoping, and summarisation.

Top performing results are currently obtained using convolutional networks (convnets) [24, 6, 25, 3, 18, 36]. Like most deep learning techniques, convnets for pixel-level object tracking benefit from large amounts of training data. Current state-of-the-art methods rely, for instance, on pixel accurate foreground/background annotations of $\sim2k$ video frames [24, 6] or $\sim10k$ images [25]. Labelling videos at the pixel level is a laborious task (compared e.g. to drawing bounding boxes for detection), and creating a large training set requires significant annotation effort.

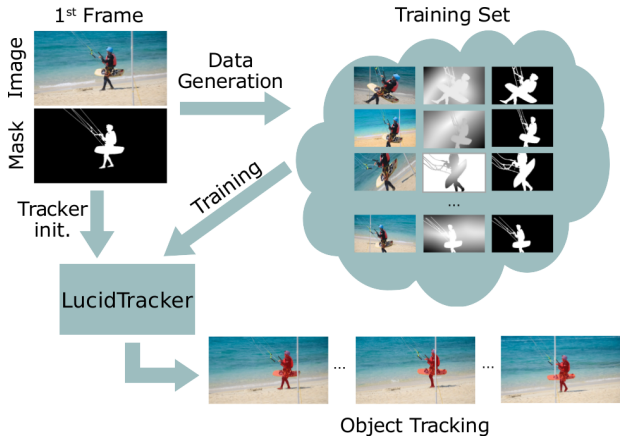

Figure 1: Starting from scarce annotations we synthesize in-domain data to train a specialized pixel-level object tracker for each dataset or even each video.

In this work we aim to reduce the necessity for such large volumes of training data. It is traditionally assumed that convnets requires large training sets to perform best. We show that for video object tracking having a larger training set is not automatically better and that improved results can be obtained by using $20\times\sim100\times$ less training data than previous approaches [6, 25]. The main insight of our work is that for pixel-level object tracking using few training frames ($1\sim100$) in the target domain is more useful than using large training volumes across domains ($1k\sim10k$).

To ensure a sufficient amount of training data close to the target domain, we develop a new technique for synthesizing training data particularly tailored for the object tracking scenario. We call this data generation strategy "*lucid dreaming*", where the first frame and its annotation mask are used to generate plausible future frames of the videos. The goal is to produce a large training set of reasonably realistic images which capture the expected appearance variations in future video frames, and thus is, by design, close to the

[1]In a lucid dream the sleeper is aware that he or she is dreaming and is sometimes able to control the course of the dream.

target domain.

Enabled by the proposed data generation strategy and the efficient use of optical flow, we are able to achieve high quality results while using only $\sim$100 individual annotated training frames. Moreover, in the extreme case with only a single annotated frame (zero pre-training), we still obtain competitive tracking results.

In summary, our contributions are the following:

1. We propose "lucid data dreaming", an automated approach to synthesize training data for the convnet-based pixel-level object tracking that enables to reach the state-of-the-art results on three different datasets.
2. We conduct an extensive analysis to explore the factors contributing to our good results.
3. We show that training a convnet for object tracking can be done with only few annotated frames. We hope these results will affect the trend towards even larger training sets, and popularize the design of trackers with lighter training needs.

2. Related work

Box-level tracking. Classic work on video object tracking focused on bounding box tracking. Many of the insights from these works have been re-used for pixel-level tracking. Traditional box tracking smoothly updates across time a linear model over hand-crafted features [19, 5, 29]. Since then, convnets have been used as improved features [11, 32, 58], and eventually to drive the tracking itself [18, 3, 53, 35, 36]. Convnet-based approaches need data for pre-training and learning the task.

Pixel-level tracking. In this paper we focus on generating a foreground versus background pixel-wise object labelling for each video frame starting from a first manually annotated frame. Multiple strategies have been proposed to solve this task.

Box-to-segment: First a box-level track is built, and a space-time grabcut-like approach is used to generate per frame segments [63].

Video saliency: Instead of tracking, these methods extract the main foreground object pixel-level space-time tube. Both hand-crafted models [13, 38] or trained convnets [54, 23] have been considered. Because these methods ignore the first frame annotation, they fail in videos where multiple salient objects move (e.g. flock of penguins).

Video segmentation methods partition the space-time volume, and then the tube overlapping most with the first frame annotation is selected as tracking output [16, 42, 7].

Mask propagation: Appearance similarity and motion smoothness across time is used to propagate the first frame annotation across the video [33, 60, 55]. These methods usually leverage optical flow and long term trajectories.

Convnets: following the trend in box-level tracking, recently convnets have been proposed for pixel-level tracking. [6] trains a generic object saliency network, and fine-tunes it per-video (using the first frame annotation) to make the output sensitive to the specific object instance being tracked. [25] uses a similar strategy, but also feeds the mask from the previous frame as guidance for the saliency network. Finally [24] mixes convnets with ideas of bilateral filtering. Our approach belongs to this group.

What makes convnets particularly suitable for the task, is that they can learn what are the common statistics of appearance and motion patterns of objects, as well as what makes them distinctive from the background, and exploit this knowledge when tracking a single particular object. This aspect gives convnets an edge over traditional techniques based on low-level features.

Our network architecture is similar to [6, 25]. Other than implementation details, there are two differentiating factors. One, we use a different strategy for training: [6, 24] rely on consecutive video training frames and [25] uses an external saliency dataset, while our approach focuses on using the first frame annotations provided with each targeted video benchmark without relying on external annotations. Two, our approach exploits optical flow better than these previous methods.

Interactive video segmentation. Interactive segmentation [34, 22, 48, 59] uses more flexible user inputs (e.g. strokes), and requires interactive processing speed rather than providing maximal quality. Albeit our technique can be adapted for more flexible inputs, we focus on maximizing quality for the non-interactive case with no-additional hints along the video.

Semantic labelling. Like other convnets in this space [24, 6, 25], our architecture builds upon the insights from the semantic labelling networks [65, 31, 62, 2]. Because of this, the flurry of recent developments should directly translate into better tracking results. For the sake of comparison with previous work, we build upon the well established VGG DeepLab architecture [8].

Synthetic data. Like in our approach, previous works have also explored synthesizing training data. Synthetic renderings [37], video game environment [45], mix-synthetic and real images [56, 9, 12] have shown promise, but require task-appropriate 3d models. Compositing real images provides more realistic results, and has shown promise for object detection [14, 52], text localization [17] and pose estimation [43].

The closest work to ours is [39], which also generates video-specific training data using the first frame annotations. They use human skeleton annotations to improve pose estimation, while we employ mask annotations to improve object tracking.

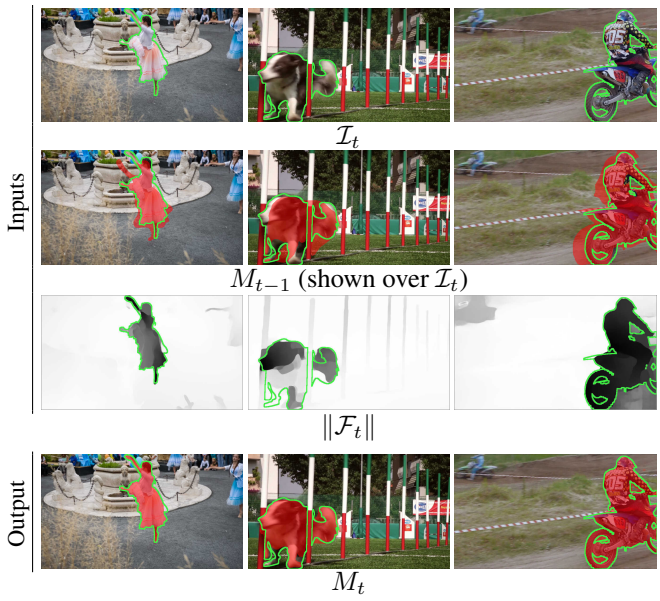

Figure 2: Data flow examples. $\mathcal{I}_t$, $\mathcal{F}_t$, M_{t-1} are the inputs, M_t is the resulting output. Green boundaries outline the ground truth segments. Red overlay indicates M_{t-1}, M_t.

3. LucidTracker

Section 3.1 describes the network architecture used, and how RGB and optical flow information are fused to predict the next frame segmentation mask. Section 3.2 discusses different training modalities employed with the proposed object tracking system.

3.1. Architecture

Approach. We model the pixel-level object tracking problem as a mask refinement task (mask: binary foreground/background labelling of the image) based on appearance and motion cues. From frame $t-1$ to frame t the estimated mask M_{t-1} is propagated to frame t, and the new mask M_t is computed as a function of the previous mask, the new image $\mathcal{I}_t$, and the optical flow $\mathcal{F}_t$, i.e. $M_t = f(\mathcal{I}_t, \mathcal{F}_t, M_{t-1})$. Since objects have a tendency to move smoothly through space in time, there are little changes from frame to frame and mask M_{t-1} can be seen as a rough estimate of M_t. Thus we require our trained convnet to learn to refine rough masks into accurate masks. Fusing the complementary image $\mathcal{I}_t$ and motion $\mathcal{F}_t$ cues exploits the information inherent to video and enables the model to segment well both static and moving objects.

Note that this approach is incremental, does a single forward pass over the video, and keeps no explicit model of the object appearance at frame t. In some experiments we adapt the model f per video, using the annotated first frame $\mathcal{I}_0$, M_0. However, in contrast to traditional techniques [19], this model is not updated while we process the video frames, thus the only state evolving along the video is the mask M_{t-1} itself.

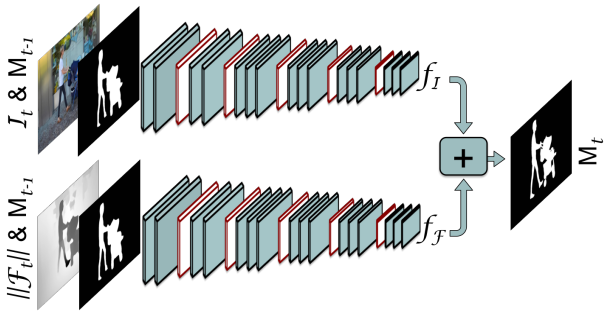

Figure 3: Overview of our architecture, where image $\mathcal{I}_t$ and optical flow information $\|\mathcal{F}_t\|$ are used to update mask M_{t-1} into M_t. See equation 1 and §3.1.

First frame. In the video object tracking task the mask for the first frame M_0 is given. This is the standard protocol of the benchmarks considered in Section 5. If only a bounding box is available on the first frame, then the mask could be estimated using grabcut-like techniques [46, 51].

RGB image $\mathcal{I}$. Typically a semantic labeller generates pixel-wise labels based on the input image (e.g. $M = g(\mathcal{I})$). We use an augmented semantic labeller with an input layer modified to accept 4 channels (RGB + previous mask) so as to generate outputs based on the previous mask estimate, e.g. $M_t = f_{\mathcal{I}}(\mathcal{I}_t, M_{t-1})$. Our approach is general and can leverage any existing semantic labelling architecture. We select the DeepLabv2 architecture with VGG base network [8], which is comparable to [24, 6, 25]; FusionSeg [23] uses ResNet.

Optical flow $\mathcal{F}$. We use flow in two complementary ways. First, to obtain a better estimate of M_t we warp M_{t-1} using the flow $\mathcal{F}_t$: $M_t = f_{\mathcal{I}}(\mathcal{I}_t, w(M_{t-1}, \mathcal{F}_t))$. Second, we use flow as a direct source of information about the mask M_t. As can be seen in Figure 2, when the object is moving relative to background, the flow magnitude $\|\mathcal{F}_t\|$ provides a very reasonable estimate of the mask M_t. We thus consider using convnet specifically for mask estimation from flow: $M_t = f_{\mathcal{F}}(\mathcal{F}_t, w(M_{t-1}, \mathcal{F}_t))$, and merge it with the image-only version by naive averaging

$$M_t = 0.5 \cdot f_{\mathcal{I}}(\mathcal{I}_t, \ldots) + 0.5 \cdot f_{\mathcal{F}}(\mathcal{F}_t, \ldots). \quad (1)$$

We use the state-of-the-art optical flow method FlowNet2.0 [20], which itself is a convnet that computes $\mathcal{F}_t = h(\mathcal{I}_{t-1}, \mathcal{I}_t)$ and is trained on synthetic renderings of flying objects [37].

The loss function is the sum of cross-entropy terms over each pixel in the output map (all pixels are equally weighted). In our experiments $f_{\mathcal{I}}$ and $f_{\mathcal{F}}$ are trained independently, via some of the modalities listed in 3.2.

We also explored expanding our network to accept 5 input channels (RGB + previous mask + flow magnitude):

$M_t = f_{\mathcal{I}+\mathcal{F}}(\mathcal{I}_t, \mathcal{F}_t, w(M_{t-1}, \mathcal{F}_t))$, but did not observe improved results compared to naive averaging. Our two streams architecture is illustrated in Figure 3.

Post-processing. As a final stage of our pipeline, we refine the generated mask M_t using DenseCRF [26] per frame. This adjusts small image details that the network might not have captured. It is known by practitioners that DenseCRF is quite sensitive to its parameters and can easily worsen results. We will use our lucid dreams to handle per-dataset CRF-tuning too, see Section 3.2.

We refer to our full system as `LucidTracker`, and as `LucidTracker`$^-$ when no post-processing is used.

3.2. Training modalities

Multiple modalities are available to train a tracker. **Training-free** approaches (e.g. BVS [33], SVT [60]) are fully hand-crafted systems with hand-tuned parameters, and thus do not require training data. They can be used as-is over different datasets. Supervised methods can also be trained to generate a **dataset-agnostic** model that can be applied over different datasets. Instead of using a fixed model for all cases, it is also possible to obtain specialized **per-dataset** models, either via self-supervision [61, 40, 64, 66] or by using the first frame annotation of each video in the dataset as training/tuning set. Finally, inspired by traditional tracking techniques, we also consider adapting the model weights to the specific video at hand, thus obtaining **per-video** models. Section 5 reports new results over these four training modalities (training-free, dataset-agnostic, per-dataset, and per-video).

Our `LucidTracker` obtains best results when first pre-trained on ImageNet, then trained per-dataset using all data from first frame annotations together, and finally fine-tuned per-video for each evaluated sequence. The post-processing DenseCRF stage is tuned per-dataset (parameter grid search). The experimental section 5 details the effect of these training stages; we can obtain reasonable performance even when training from only a single annotated frame (without ImageNet pre-training).

Unless otherwise stated, we fine-tune per-video models relying solely on the first frame $\mathcal{I}_0$ and its annotation M_0. This is in contrast to traditional techniques [19, 5, 29] which would update the appearance model at each frame.

4. Lucid data dreaming

To train the function f one would think of using ground truth data for M_{t-1} and M_t (like [3, 6, 18]), however such data is expensive to annotate and rare. [6] thus trains on a set of 30 videos ($\sim 2k$ frames) and requires the model to transfer across multiple tests sets. [25] side-steps the need for consecutive frames by generating synthetic masks M_{t-1} from a large saliency dataset of $\sim 10k$ images with their corresponding mask M_t . We propose a new data generation strategy to reach better results using only ~ 100 individual training frames.

Ideally training data should be as similar as possible to the test data, even subtle differences may affect quality (e.g. training on static images for testing on videos underperforms [50]). To ensure our training data is in-domain, we propose to generate it by synthesizing samples from the provided annotated frame (first frame) in each target video. This is akin to "lucid dreaming" as we intentionally "dream" the desired data, by creating sample images that are plausible hypothetical future frames of the video. The outcome of this process is a large set ($2.5k$ pairs per annotation) of frame pairs in the target domain with known optical flow and mask annotations, see Figure 4.

Synthesis process. The target domain for a tracker is the set of future frames of the given video. Traditional data augmentation via small image perturbation is insufficient to cover the expect variations across time, thus a task specific strategy is needed. Across the video the tracked object might change in illumination, deform, translate, be occluded, show different point of views, and evolve on top of a dynamic background. All of these aspects need to be captured when synthesizing future frames. We achieve this by cutting-out the foreground object, in-painting the background, perturbing both foreground and background, and finally recomposing the scene. This process is applied twice with randomly sampled transform parameters, resulting in a pair of frames ($\mathcal{I}_{\tau-1}$, $\mathcal{I}_\tau$) with ground-truth pixel-level mask annotations ($M_{\tau-1}$, M_τ), optical flow $\mathcal{F}_\tau$, and occlusion regions, as the undergoing transformations are known. The object position in $\mathcal{I}_\tau$ is uniformly sampled, but the changes between $\mathcal{I}_{\tau-1}$, $\mathcal{I}_\tau$ are kept small to mimic the usual evolution between consecutive frames.
In more details, starting from an annotated image:
1. Illumination changes: we globally modify the image by randomly altering saturation S and value V (from HSV colour space) via $x' = a \cdot x^b + c$, where $a \in 1 \pm 0.05$, $b \in 1 \pm 0.3$, and $c \in \pm 0.07$.
2. Fg/Bg split: the foreground object is removed from the image $\mathcal{I}_0$ and a background image is created by inpainting the cut-out area [10].
3. Object motion: we simulate motion and shape deformations by applying global translation as well as affine and non-rigid deformations to the foreground object. For $\mathcal{I}_{\tau-1}$ the object is placed at any location within the image with a uniform distribution, and in $\mathcal{I}_\tau$ with a translation of $\pm 10\%$ of the object size relative to $\tau - 1$. In both frames we apply random rotation $\pm 30°$, scaling $\pm 15\%$ and thin-plate splines deformations [4] of $\pm 10\%$ of the object size.
4. Camera motion: We additionally transform the background using affine deformations to simulate camera view changes. We apply here random translation, rotation, and

(a) Original image $\mathcal{I}_0$ and mask annotation M_0

(b) Generated image $\mathcal{I}_{\tau-1}$

(c) Generated image $\mathcal{I}_\tau$

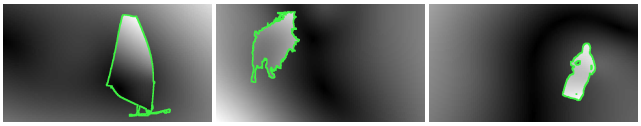

(d) Generated flow magnitude $\|\mathcal{F}_\tau\|$

Figure 4: Lucid data dreaming examples. From one annotated frame we generate pairs of images $(\mathcal{I}_{\tau-1}, \mathcal{I}_\tau)$ that are plausible future video frames, with known optical flow $(\mathcal{F}_\tau)$ and masks (green boundaries). Note the inpainted background and foreground/background deformations.

scaling within the same ranges as for the foreground object.
5. Fg/Bg merge: finally $(\mathcal{I}_{\tau-1}, \mathcal{I}_\tau)$ are composed by blending the perturbed foreground with the perturbed background using Poisson matting [49]. Using the known transformation parameters we also synthesize ground-truth pixel-level mask annotations $(M_{\tau-1}, M_\tau)$ and optical flow $\mathcal{F}_\tau$.
Figure 4 shows example results. Albeit our approach does not capture appearance changes due to point of view, occlusions, nor shadows, we see that already this rough modelling is effective to train our tracking models.

The number of synthesized images can be arbitrarily large. We generate $2.5k$ pairs per annotated video frame. This training data is, by design, in-domain with regard of the target video. The experimental section 5 shows that this strategy is more effective than using thousands of manually annotated images from close-by domains.

5. Results

We present here a detailed empirical evaluation on three different datasets.

5.1. Experimental setup

Datasets. We evaluate our method on three video object segmentation datasets: DAVIS [41], YouTubeObjects [44, 21], and SegTrack$_{v2}$ [30]. The goal is to track an object through all video frames given a foreground object mask in the first frame. These three datasets provide diverse challenges with a mix of HD and low-res web videos, single or multiple salient objects per video, videos with flocks of similar looking instances, longer (~ 400) and shorter (~ 10) sequences, as well as the usual tracking challenges such as occlusion, fast motion, illumination, view point changes, elastic deformation, etc. DAVIS, YouTubeObjects, and SegTrack$_{v2}$ contain respectively 50, 126, and 14 videos covering 50, 126, 24 object instances each. Each instance is treated as a separate problem.

We use DAVIS as main experimental playground, and report on two other datasets as complementary test set results.

Evaluation metric. To measure the accuracy of video object tracking we use the mean intersection-over-union overlap (mIoU) between the ground truth object mask and the predicted segmentation, averaged across all video sequences. We have noticed disparate evaluation procedures used in previous work, and we report here a unified evaluation across datasets. When possible, we re-evaluated certain methods using results provided by their authors. For all three datasets we follow the DAVIS evaluation protocol, excluding the first frame from evaluation and using all other frames from the video sequences, independent of object presence in the frame. Additional details of the datasets and the evaluation are in the supplementary material.

Training details. For training all the models we use SGD with mini-batches of 10 images and a fixed learning policy with initial learning rate of 10^{-3}. The momentum and weight decay are set to 0.9 and $5 \cdot 10^{-4}$, respectively.

Models using pre-training are initialized with weights trained for image classification on ImageNet [47]. We then train per-dataset for 40k iterations with the RGB+Mask branch $f_\mathcal{I}$ and for 20k iterations for the Flow+Mask $f_\mathcal{F}$.

Models without ImageNet pre-training are initialized using the "Xavier" strategy [15]. The per-dataset training needs to be longer, using 100k iterations for the $f_\mathcal{I}$ branch and 40k iterations for the $f_\mathcal{F}$ branch.

For per-video fine-tuning 2k iterations are used for $f_\mathcal{I}$.

All training parameters are chosen based on DAVIS results. We use identical parameters on YouTubeObjects and SegTrack$_{v2}$, showing the generalization of our approach.

It takes ~3.5h to obtain each per-video model, including data generation, per-dataset training, per-video fine-tuning and per-dataset grid search of CRF parameters (averaged over DAVIS, amortizing the training time over all videos). At test time our LucidTracker runs at ~5s per frame, including the optical flow computation with FlowNet2.0 [20] (~0.5s) and CRF post-processing [26] (~2s).

5.2. Key results

Table 1 presents our main result and compares it to previous work. Our full system, `LucidTracker`, provides the best tracking quality across three datasets while being trained on each dataset using only one frame per video

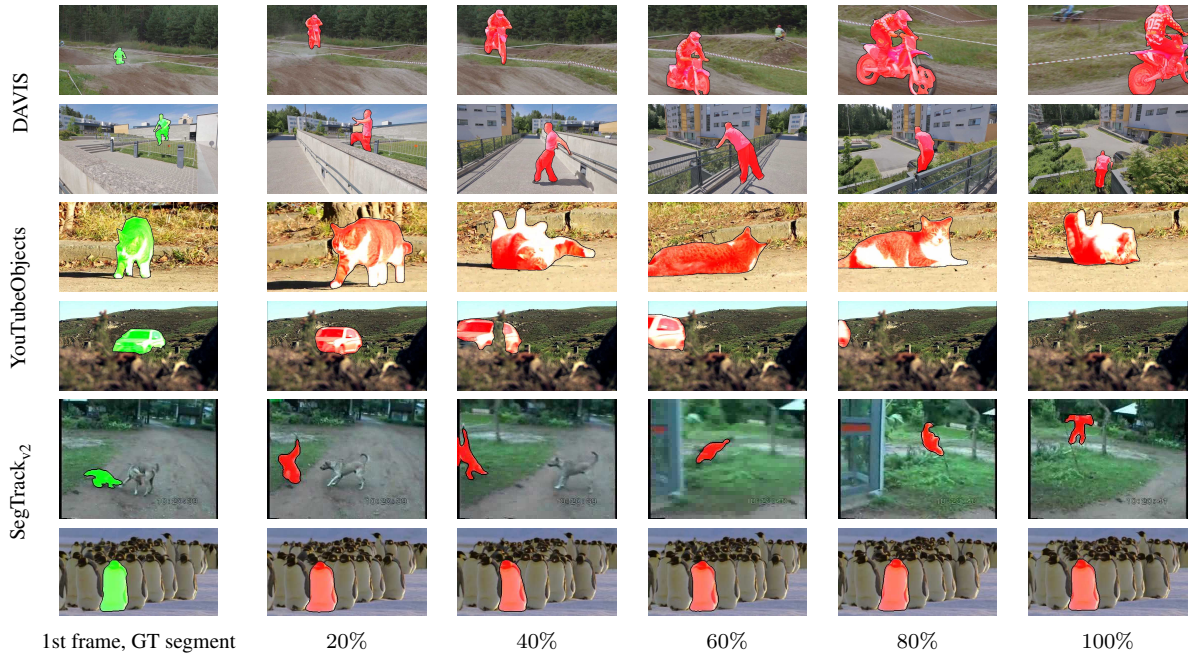

Figure 5: LucidTracker qualitative results. Frames sampled along the video duration (e.g. 50%: video middle point). Our model is robust to various challenges, such as view changes, fast motion, shape deformations, and out-of-view scenarios.

(50 frames for DAVIS, 126 for YouTubeObjects, 24 for SegTrack$_{v2}$), which is $20\times \sim 100\times$ less than the top competing methods. Ours is the first method to reach $>$ 75 mIoU on all three datasets.

Oracles and baselines. Grabcut oracle computes grabcut [46] using the ground truth bounding boxes (box oracle). This oracle indicates that on the considered datasets separating foreground from background is not easy, even if a perfect box-level tracker was available.

We provide three additional baselines. "Saliency" corresponds to using the generic (training-free) saliency method EQCut [1] over the RGB image $\mathcal{I}_t$. "Flow saliency" does the same, but over the optical flow magnitude $\|\mathcal{F}_t\|$. Results indicate that the objects being tracked are not particularly salient in the image. On DAVIS motion saliency is a strong signal but not on the other two datasets. Saliency methods ignore the first frame annotation provided for the tracking task. We also consider the "Mask warping" baseline which uses optical flow to propagate the mask estimate from t to $t+1$ via simple warping $M_t = w(M_{t-1}, \mathcal{F}_t)$. The bad results of this baseline indicate that the high quality flow [20] that we use is by itself insufficient to solve the tracking task, and that indeed our proposed convnet does the heavy lifting.

The large fluctuation of the relative baseline results across the three datasets empirically confirms that each of them presents unique challenges.

Comparison. Compared to flow propagation methods such as BVS, N15, ObjFlow, and STV, we obtain better results because we build per-video a stronger appearance model of the tracked object (embodied in the fine-tuned model). Compared to convnet learning methods such as VPN, OSVOS, MaskTrack, we require significantly less training data, yet obtain better results.

	Method	# training images	Flow $\mathcal{F}$	Dataset, mIoU DAVIS	YoutbObjs	SegTrck$_{v2}$
	Box oracle [25]	0	✗	45.1	55.3	56.1
	Grabcut oracle [25]	0	✗	67.3	67.6	74.2
Ignores 1st frame annotation	Saliency	0	✗	32.7	40.7	22.2
	NLC [13]	0	✓	64.1	-	-
	TRS [63]	0	✓	-	-	69.1
	MP-Net [54]	~22.5k	✓	69.7	-	-
	Flow saliency	0	✓	70.7	36.3	35.9
	FusionSeg [23]	~95k	✓	71.5	67.9	-
Uses 1st frame annotation	Mask warping	0	✓	32.1	43.2	42.0
	FCP [42]	0	✓	63.1	-	-
	BVS [33]	0	✗	66.5	59.7	58.4
	N15 [34]	0	✓	-	-	69.6
	ObjFlow [55]	0	✓	71.1	70.1	67.5
	STV [60]	0	✓	73.6	-	-
	VPN [24]	~2.3k	✗	*75.0*	-	-
	OSVOS [6]	~2.3k	✗	*79.8*	72.5	65.4
	MaskTrack [25]	~11k	✓	80.3	72.6	70.3
	LucidTracker	**24~126**	✓	**84.8**	**76.2**	**77.6**

Table 1: Comparison of segment tracking results across three datasets. Numbers in italic are reported on subsets of DAVIS. Our LucidTracker consistently improves over previous results, see §5.2.

Figure 5 provides qualitative results of `LucidTracker` across three different datasets. Our system is robust to vari-

Variant	$\mathcal{I}$	$\mathcal{F}$	warp. w	Dataset, mIoU DAVIS	YoutbObjs	SegTrck$_{v2}$
LucidTracker	✓	✓	✓	**84.8**	**76.2**	**77.6**
LucidTracker$^-$	✓	✓	✓	83.7	76.2	76.8
No warping	✓	✓	✗	82.0	74.6	70.5
No OF	✓	✗	✗	78.0	74.7	61.8
OF only	✗	✓	✓	74.5	43.1	55.8

Table 2: Ablation study of flow ingredients. Flow complements image only results, with large fluctuations across datasets. See §5.3.1.

ous challenges present in videos. It handles well camera view changes, fast motion, object shape deformation, out-of-view scenarios, multiple similar looking objects and even low quality video data as for the "monkeydog" video sequence in SegTrack$_{v2}$. On DAVIS `LucidTracker` is best on 13 out of 15 video attribute categories. More quantitative and qualitative results, and discussion of failure cases are in the supplementary material.

Conclusion. We show that using less training data, does not necessarily lead to poorer results. We report the best known results for this task while using only $24 \sim 126$ training frames.

5.3. Ablation studies

In this section we explore in more details how the different ingredients contribute to our results.

5.3.1 Effect of optical flow

Table 2 shows the effect of using optical flow on `LucidTracker` results. Comparing our full system to the "No OF" row, we see that the effect of optical flow varies across datasets, from minor improvement in YouTubeObjects, to major difference in SegTrack$_{v2}$. In this last dataset, using mask warping is particularly useful too.

OSVOS [6] also does not use optical flow, but instead uses a per-frame mask post-processing based on a boundary detector (trained on further external data), which provides ~ 2 percent point gain. Accounting for this, our "No OF" (and no CRF) result matches theirs on DAVIS and YouTubeObjects despite using significantly less training data (see Table 1, e.g. $79.8 - 2 \approx 78.0$ on DAVIS).

Conclusion. The results show that, as expected, flow provides a complementary signal to RGB image only. Despite its simplicity our fusion strategy ($f_{\mathcal{I}} + f_{\mathcal{F}}$) provides gains on all datasets, and leads to competitive results.

5.3.2 Effect of training modalities

Table 3 compares the effect of different ingredients in the `LucidTracker`$^-$ training. Results are obtained using RGB and flow, with warping, no CRF; $M_t = f\,(\mathcal{I}_t, w(M_{t-1}, \mathcal{F}_t))$.

Variant	ImgNet pre-train.	per-dataset training	per-video fine-tun.	Dataset, mIoU DAVIS	YoutbObjs	SegTrck$_{v2}$
LucidTracker$^-$	✓	✓	✓	**83.7**	**76.2**	**76.8**
(no ImgNet)	✗	✓	✓	82.0	74.3	71.2
No per-video	✓	✓	✗	82.7	72.3	71.9
tuning	✗	✓	✗	78.4	69.7	68.2
Only per-	✓	✗	✓	79.4	-	70.4
-video tuning	✗	✗	✓	*80.5*	-	*66.8*

Table 3: Ablation study of training modalities. ImageNet pre-training and per-video tuning provide additional improvement over per-dataset training. Even with one frame annotation for only per-video tuning we obtain good performance. See §5.3.2.

We see that ImageNet pre-training does provide $2 \sim 5$ percent point improvement (depending on the dataset of interest; e.g. $82.0 \rightarrow 83.7$ mIoU on DAVIS). Per-video fine-tuning (after doing per-dataset training) provides an additional $1 \sim 2$ percent point gain (e.g. $82.7 \rightarrow 83.7$ mIoU on DAVIS). Both ingredients clearly contribute to the tracking results.

In the bottom row ("only per-video tuning"), the model is trained per-video without ImageNet pre-training nor per-dataset training, i.e. using a *single annotated training frame*. Our network is based on VGG16 [8] and contains $\sim 20M$ parameters, all effectively learnt from a single annotated image that is augmented to become $2.5k$ training samples (see Section 4). Even with such minimal amount of training data, we still obtain a surprisingly good performance (compare 80.5 on DAVIS to others in Table 1). This shows how effective is, by itself, the proposed training strategy based on lucid dreaming of the data.

Note that training a model using only per-video tuning takes about one full GPU day per video sequence; making these results insightful but not decidedly practical.

Preliminary experiments evaluating on DAVIS the impact of the different ingredients of our lucid dreaming data generation showed, depending on the exact setup, $3 \sim 10$ percent mIoU points fluctuations between a basic version (e.g. without non-rigid deformations nor scene re-composition) and the full synthesis process described in Section 4. Having a sophisticated data generation process directly impacts the tracking quality.

Conclusion. Both ImageNet pre-training and per-video tuning of the models provide complementary gains over the default per-dataset training. Per-video training by itself, despite using a single annotated frame, provides already much of the needed information for the tracking task.

5.4. Additional experiments

Other than adding or removing ingredients, as in Section 5.3, we also want to understand how the training data itself affects the obtained results.

	# training videos	# frames per video	mIoU
Uses 1st frames from test set	1	1	78.3
	2	1	75.4
	15	1	68.7
	30	1	65.4
	30	2	74.3
Ignores 1st frames from test set	2	1	11.6
	15	1	36.4
	30	1	41.7
	30	2	48.4

Table 4: Varying the number of training videos. A smaller training set closer to the target domain is better than a larger one. See §5.4.1.

5.4.1 Generalization across videos

Table 4 explores the effect of tracking quality as a function of the number of training samples. It reports the results with varying the number of training videos and with/without including the first frames of each test video in the training. To see more directly the training data effects we use a base model with RGB image $\mathcal{I}_t$ only (no flow $\mathcal{F}$, no CRF), and per-dataset training (no ImageNet pre-training). We evaluate on two subsets of 15 DAVIS videos each, where the first frames for training are taken from only one subset. The reported numbers are thus comparable within Table 4, but not across to the other tables in the paper.

Comparing the top and bottom sections, we see that when the annotations from the test set videos are not included, tracking quality drops drastically (e.g. $68.7 \rightarrow 36.4$ mIoU). Conversely, on subset of videos for which the first frame annotation is used for training, the quality is much higher and improves as the training samples become more and more specific (in-domain) to the target video ($65.4 \rightarrow 78.3$ mIoU). Adding extra videos for training does not improve the performance. It is better ($68.7 \rightarrow 78.3$ mIoU) to have 15 models each trained and evaluated on a single video (row 1-1-✓) than having one model trained over 15 test videos (row 15-1-✓).

Training with an additional frame from each video (we added the last frame of each train video) significantly boosts the resulting within-video quality ($65.4 \rightarrow 74.3$ mIoU), because the training samples cover better the test domain.

Conclusion. These results show that, when using RGB information ($\mathcal{I}_t$), increasing the number of training videos *does not* improve the resulting quality of our system. Even within a dataset, properly using the training sample(s) from within each video matters more than collecting more videos to build a larger training set.

5.4.2 Generalization across datasets

Section 5.4.1 has explored the effect of changing the volume of training data within one dataset, Table 5 compares results when using different datasets for training. Results are obtained using a base model with RGB and flow ($M_t = f(\mathcal{I}_t, M_{t-1})$, no warping, no CRF), ImageNet pre-training, per-dataset training, and no per-video tuning to accentuate the effect of the training dataset.

Training set	Dataset, mIoU DAVIS	YoutbObjs	$\text{SegTrck}_{\text{v2}}$	Mean
DAVIS	80.9	50.9	46.9	59.6
YoutbObjs	*67.0*	71.5	*52.0*	63.5
$\text{SegTrack}_{\text{v2}}$	56.0	*52.2*	66.4	58.2
Best	**80.9**	**71.5**	**66.4**	**72.9**
Second best	67.0	52.2	52.0	57.1
All-in-one	71.9	70.7	60.8	67.8

Table 5: Generalization across datasets. Results with underline are the best per dataset, and in italic are the second best per dataset (ignoring all-in-one setup). We observe a significant quality gap between training from the target videos, versus training from other datasets; see §5.4.2.

The best performance is obtained when training on the first frames of the target set. There is a noticeable $\sim$10 percent points drop when moving to the second best choice (e.g. $80.9 \rightarrow 67.0$ for DAVIS). Interestingly, when putting all the datasets together for training ("all-in-one" row, a dataset-agnostic model) results degrade, reinforcing the idea that "just adding more data" does not automatically make the performance better.

Conclusion. Best results are obtained when using training data that focuses on the test video sequences, using similar datasets or combining multiple datasets degrades the performance for our system.

6. Conclusion

We have described a new convnet-based approach for pixel-level object tracking in videos. In contrast to previous work, we show that top results in three datasets can be achieved without requiring external training datasets (neither images with saliency annotation nor annotated videos). Even more, our experiments indicate that it is not always beneficial to use additional training data, synthesizing training samples close to the test domain is more effective than adding more training samples from related domains.

Our extensive analysis decomposed the ingredients that contribute to our improved results, indicating that our new training strategy and the way we handle optical flow are key.

Showing that training a convnet for object tracking can be done with only few (~ 100) training samples changes the mindset regarding how much general "objectness" knowledge is required to approach this problem [25, 23], and more broadly how much training data is required to train large convnets depending on the task at hand.

We hope these new results will fuel the ongoing evolution of convnet techniques for object tracking.

Acknowledgements

Eddy Ilg and Thomas Brox acknowledge funding by the DFG Grant BR 3815/7-1.

References

[1] Ç. Aytekin, E. C. Ozan, S. Kiranyaz, and M. Gabbouj. Visual saliency by extended quantum cuts. In *ICIP*, 2015. 6

[2] A. Bansal, X. Chen, B. Russell, A. Gupta, and D. Ramanan. Pixelnet: Representation of the pixels, by the pixels, and for the pixels. *arXiv:1702.06506*, 2017. 2

[3] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr. Fully-convolutional siamese networks for object tracking. *arXiv:1606.09549*, 2016. 1, 2, 4

[4] F. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *PAMI*, 1989. 4

[5] M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool. Robust tracking-by-detection using a detector confidence particle filter. In *ICCV*, 2009. 2, 4

[6] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixe, D. Cremers, and L. V. Gool. One-shot video object segmentation. In *CVPR*, 2017. 1, 2, 3, 4, 6, 7, 13, 14, 15

[7] J. Chang, D. Wei, and J. W. Fisher. A video representation using temporal superpixels. In *CVPR*, 2013. 2

[8] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv:1606.00915*, 2016. 2, 3, 7, 13

[9] W. Chen, H. Wang, Y. Li, H. Su, Z. Wang, C. Tu, D. Lischinski, D. Cohen-Or, and B. Chen. Synthesizing training images for boosting human 3d pose estimation. In *3D Vision (3DV)*, 2016. 2

[10] A. Criminisi, P. Perez, and K. Toyama. Region filling and object removal by exemplar-based image inpainting. *Trans. Img. Proc.*, 2004. 4

[11] M. Danelljan, G. Hager, F. Shahbaz Khan, and M. Felsberg. Convolutional features for correlation filter based visual tracking. In *ICCV workshop*, 2015. 2

[12] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazirbas, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. Flownet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 2

[13] A. Faktor and M. Irani. Video segmentation by non-local consensus voting. In *BMVC*, 2014. 2, 6, 14

[14] G. Georgakis, A. Mousavian, A. C. Berg, and J. Kosecka. Synthesizing training data for object detection in indoor scenes. *arXiv:1702.07836*, 2017. 2

[15] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. 5

[16] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *CVPR*, 2010. 2

[17] A. Gupta, A. Vedaldi, and A. Zisserman. Synthetic data for text localisation in natural images. In *CVPR*, 2016. 2

[18] D. Held, S. Thrun, and S. Savarese. Learning to track at 100 fps with deep regression networks. In *ECCV*, 2016. 1, 2, 4

[19] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. Exploiting the circulant structure of tracking-by-detection with kernels. In *ECCV*, 2012. 2, 3, 4

[20] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 3, 5, 6

[21] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In *ECCV*, 2014. 5, 11

[22] S. D. Jain and K. Grauman. Click carving: Segmenting objects in video with point clicks. In *HCOMP*, 2016. 2

[23] S. D. Jain, B. Xiong, and K. Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmention of generic objects in videos. *arXiv:1701.05384*, 2017. 2, 3, 6, 8, 14, 15

[24] V. Jampani, R. Gadde, and P. V. Gehler. Video propagation networks. *arXiv:1612.05478*, 2016. 1, 2, 3, 6, 14

[25] A. Khoreva, F. Perazzi, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *arXiv:1612.02646*, 2016. 1, 2, 3, 4, 6, 8, 13, 14, 15

[26] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*. 2011. 4, 5, 11

[27] M. Kristan, J. Matas, et al. The visual object tracking vot2015 challenge results. In *ICCV workshop*, 2015. 1

[28] M. Kristan, J. Matas, et al. The visual object tracking vot2016 challenge results. In *ECCV workshop*, 2016. 1

[29] M. Kristan, R. Pflugfelder, et al. The visual object tracking vot2014 challenge results. In *ECCV workshop*, 2014. 1, 2, 4

[30] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 1, 5, 11

[31] G. Lin, A. Milan, C. Shen, and I. D. Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. *arXiv:1611.06612*, 2016. 2

[32] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang. Hierarchical convolutional features for visual tracking. In *ICCV*, 2015. 2

[33] N. Maerki, F. Perazzi, O. Wang, and A. Sorkine-Hornung. Bilateral space video segmentation. In *CVPR*, 2016. 2, 4, 6, 14, 15

[34] N. Nagaraja, F. Schmidt, and T. Brox. Video segmentation with just a few strokes. In *ICCV*, 2015. 2, 6, 15

[35] H. Nam, M. Baek, and B. Han. Modeling and propagating cnns in a tree structure for visual tracking. *arXiv:1608.07242*, 2016. 2

[36] H. Nam and B. Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 1, 2

[37] N.Mayer, E.Ilg, P.Häusser, P.Fischer, D.Cremers, A.Dosovitskiy, and T.Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016. 2, 3

[38] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013. 2

[39] D. Park and D. Ramanan. Articulated pose estimation with tiny synthetic videos. In *CVPR Workshop*, 2015. 2

[40] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan. Learning features by watching objects move. *arXiv:1612.06370*, 2016. 4

[41] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 1, 5, 11, 13

[42] F. Perazzi, O. Wang, M. Gross, and A. Sorkine-Hornung. Fully connected object proposals for video segmentation. In *ICCV*, 2015. 2, 6, 14

[43] L. Pishchulin, A. Jain, M. Andriluka, T. Thormählen, and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR*, 2012. 2

[44] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*, 2012. 1, 5, 11

[45] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. 2

[46] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *SIGGRAPH*, 2004. 3, 6

[47] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5

[48] T. V. Spina and A. X. Falcão. Fomtrace: Interactive video segmentation by image graphs and fuzzy object models. *arXiv:1606.03369*, 2016. 2

[49] J. Sun, J. Jia, C.-K. Tang, and H.-Y. Shum. Poisson matting. In *SIGGRAPH*, 2004. 5

[50] K. Tang, V. Ramanathan, L. Fei-fei, and D. Koller. Shifting weights: Adapting object detectors from image to video. In *NIPS*, 2012. 4

[51] M. Tang, D. Marin, I. Ben Ayed, and Y. Boykov. Normalized cut meets mrf. In *ECCV*, 2016. 3

[52] S. Tang, M. Andriluka, A. Milan, K. Schindler, S. Roth, and B. Schiele. Learning people detectors for tracking in crowded scenes. In *ICCV*, 2013. 2

[53] R. Tao, E. Gavves, and A. W. Smeulders. Siamese instance search for tracking. *arXiv:1605.05863*, 2016. 2

[54] P. Tokmakov, K. Alahari, and C. Schmid. Learning motion patterns in videos. *arXiv:1612.07217*, 2016. 2, 6, 14

[55] Y.-H. Tsai, M.-H. Yang, and M. J. Black. Video segmentation via object flow. In *CVPR*, 2016. 2, 6, 14, 15

[56] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid. Learning from synthetic humans. *arXiv:1701.01370*. 2

[57] T. Vojir and J. Matas. Pixel-wise object segmentations for the VOT 2016 dataset. Research report, 2017. 1

[58] L. Wang, W. Ouyang, X. Wang, and H. Lu. Visual tracking with fully convolutional networks. In *ICCV*, 2015. 2

[59] T. Wang, B. Han, and J. Collomosse. Touchcut: Fast image and video segmentation using single-touch interaction. *CVIU*, 2014. 2

[60] W. Wang and J. Shen. Super-trajectory for video segmentation. *arXiv:1702.08634*, 2017. 2, 4, 6, 14

[61] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 4

[62] Z. Wu, C. Shen, and A. van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *arXiv:1611.10080*, 2016. 2

[63] F. Xiao and Y. J. Lee. Track and segment: An iterative unsupervised approach for video object proposals. In *CVPR*, 2016. 2, 6, 15

[64] J. J. Yu, A. W. Harley, and K. G. Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. *arXiv:1608.05842*, 2016. 4

[65] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. *arXiv:1612.01105*, 2016. 2

[66] Y. Zhu, Z. Lan, S. Newsam, and A. G. Hauptmann. Guided optical flow learning. *arXiv:1702.02295*, 2017. 4

Supplementary material

A. Content

This supplementary material provides additional quantitative and qualitative results.

- Section B provides additional lucid data dreaming examples.
- Section C details the experimental setup.
- Section D shows the importance of CRF tuning over "lucid dreams".
- Section E reports the detailed DAVIS benchmark evaluations including contour accuracy, temporal stability measures, and attribute based evaluation.
- Section F provides per-sequence results for DAVIS and SegTrack$_{v2}$.
- Section G discusses the disparate evaluation procedures on YouTubeObjects.
- Section H shows additional qualitative results and failure cases.

B. Lucid data dreaming examples

In this work we propose to generate the training data by synthesizing training samples from the provided annotated frame (first frame) in each target video. We call this process lucid data dreaming. The goal is to produce a large set of reasonably realistic images which captures most of the variability expected in the future frames. We generate new training frames by cropping out the annotated foreground object, transforming both foreground and background image and then pasting back the foreground object to create a new scene configuration. Figure S1 shows examples of the data generated by lucid dreaming strategy.

Using this "lucid dreaming" strategy from one annotated frame we generate $2.5k$ image pairs $(\mathcal{I}_{\tau-1}, \mathcal{I}_{\tau})$, with optical flow magnitude $\mathcal{F}_{\tau}$ and ground truth masks $M_{\tau-1}$, M_{τ}. The number of synthesized images can be arbitrarily large. Using data "dreamed" from the first frame of a (set of) video(s) ensures that the training data is within the domain of said video(s).

C. Experimental setup details

Datasets. We evaluate our method on three video object segmentation datasets: DAVIS [41], YouTubeObjects [44, 21], and SegTrack$_{v2}$ [30], where the goal is to track an object through all video frames given a foreground object mask in the first frame. All three datasets represent different challenges, have different types of annotations and quality of the video data.

Method	CRF parameters	Dataset, mIoU		
		DAVIS	YoutbObjs	SegTrck$_{v2}$
LucidTracker^{-}	-	83.7	76.2	76.8
LucidTracker	default	84.2	75.5	72.2
LucidTracker	tuned per-dataset	**84.8**	**76.2**	**77.6**

Table S1: Effect of CRF tuning. Without per-dataset tuning DenseCRF will under-perform.

DAVIS [41] is the most recently released video segmentation benchmark, which consists of 50 full-HD videos of diverse object categories with all frames annotated with pixel-level accuracy, where one single or two connected moving objects are separated from the background. The number of frames in each video varies from 25 to 104.

YouTubeObjects [44, 21] includes web videos from 10 object categories. We use the subset of 126 video sequences with mask annotations provided by [21] for evaluation, where one single object or a group of objects of the same category are separated from the background. In contrast to DAVIS these videos have a mix of static and moving objects. The number of frames in each video ranges from 2 to 401.

SegTrack$_{v2}$[30] consists of 14 videos with multiple object annotations for each frame. For videos with multiple objects each object is treated as a separate problem, resulting in 24 sequences. The length of each video varies from 21 to 279 frames. The images in this dataset have low resolution and some compression artefacts, making it hard to track the object based on its appearance. Therefore some fluctuations in the results could be observed due to the small set of videos and their low quality.

The main experimental part is done on DAVIS, since it is the largest densely annotated dataset out of the three, and provides high quality/high resolution data. The videos for this dataset were chosen to represent diverse challenges, making it a good experimental playground.

Evaluation. To measure the accuracy of video object tracking we use the mean intersection-over-union overlap (mIoU) between the ground truth object mask and the predicted segmentation, averaged across all video sequences. For all three datasets we follow the DAVIS evaluation protocol, excluding the first frame from evaluation and using all other frames from the video sequences, independent of object presence in the frame.

Section G compares results using our evaluation versus the ad-hoc YouTubeObjects protocol.

D. Importance of CRF tuning

As a final stage of our pipeline, we refine the generated mask using DenseCRF [26] per frame. This captures small image details that the network might have missed. It is

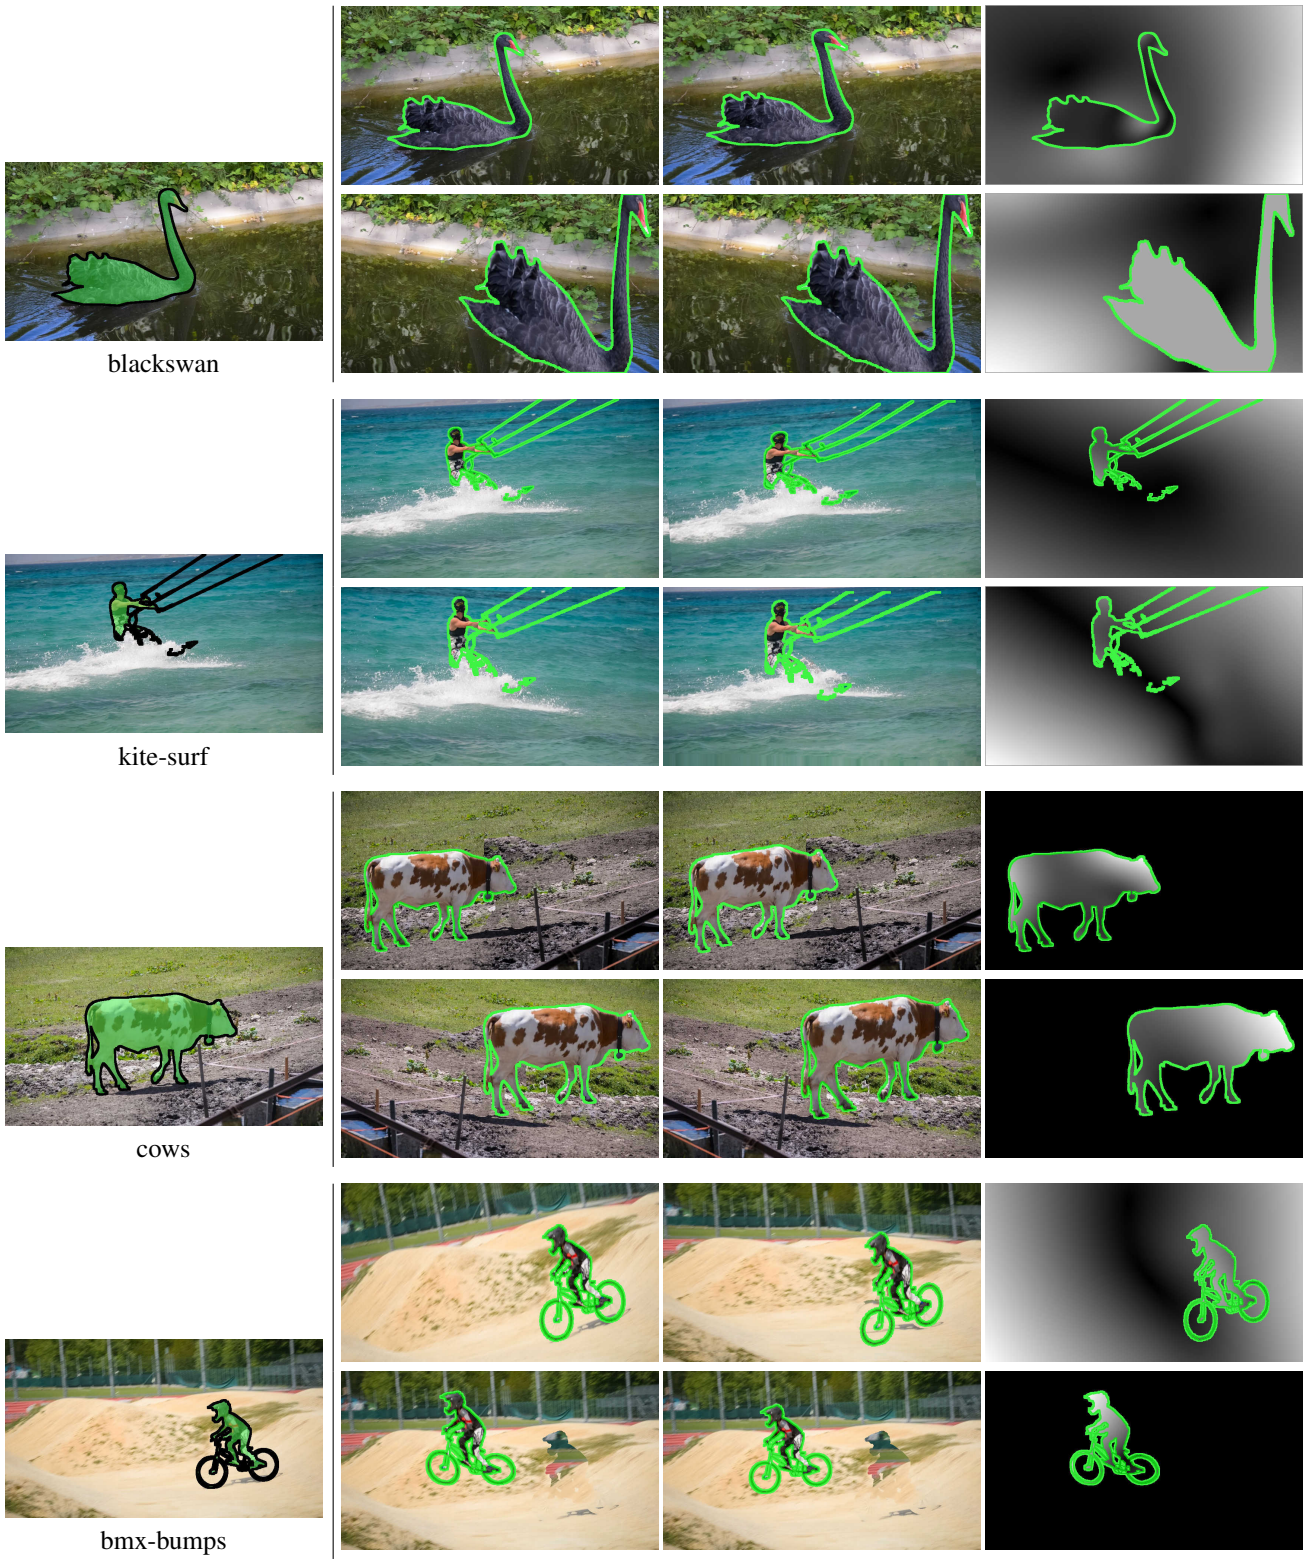

Figure S1: Lucid data dreaming examples. From one annotated frame we generate pairs of images ($\mathcal{I}_{\tau-1}$, $\mathcal{I}_\tau$) that are plausible future video frames, with known optical flow ($\mathcal{F}_\tau$) and masks (green boundaries). Note the inpainted background and foreground/background deformations.

Figure S2: Effect of CRF tunning. The videos are chosen with the highest margin between with and without CRF post-processing (based on mIoU over the video).

known by practitioners that DenseCRF is quite sensitive to its parameters and can easily worsen results. We use our lucid dreams to handle per-dataset CRF-tuning.

DenseCRF tuning details. Following [8] we employ grid search scheme for tuning CRF parameters. Once the per-dataset tracking model is trained, we apply it over a subset of its training set (5 random images from the lucid dreams), apply DenseCRF with the given parameters over this output, and then compare to the lucid dream ground truth.

The ranges of the color (RGB) and the spatial (XY) standard deviation of the DenseCRF appearance kernel are set, respectively, to [3:1:10] and [10:10:100]. The search range of the pairwise term weight is set to [3:1:5]. We employ an additional smoothness kernel to remove small isolated regions. The range for both of its weight and the spatial (XY) standard deviation is set to [1:1:3]. The number of mean field iterations is fixed to 10. We also evaluate the results when no DenseCRF is applied.

The impact of the tuned parameter of DenseCRF post-processing is shown in Table S1 and Figure S2. Table S1 indicates that using default DenseCRF parameters will degrade performance, and that per-dataset tunning is important.

E. DAVIS benchmark

Table S2 presents a more detailed evaluation on DAVIS using evaluation metrics proposed in [41]. Three measures are used: region similarity in terms of intersection over union (J), contour accuracy (F), and temporal instability of the masks (T). We outperform the competitive methods [25, 6] on all three measures.

Attribute evaluation. Table S3 reports attribute based evaluation on DAVIS. `LucidTracker` has the best performance on 13 out of 15 video attribute categories. The attribute based evaluation shows that `LucidTracker` is robust to various video challenges present in DAVIS.

F. Per-sequence results

DAVIS. We present the per-sequence comparison with other state-of-the-art methods on DAVIS in Table S4.

SegTrack$_{v2}$. Table S5 reports the per-sequence comparison with other state-of-the-art methods on SegTrack$_{v2}$.

With the exception of a couple of failure modes (1.4, 2.8 mIoU), our method consistently improves across videos on both datasets.

G. YouTubeObjects evaluation details

We have noticed disparate evaluation procedures used in previous work. On YouTubeObjects some of the evaluations excluded frames where the object is no longer present. We consider this case one of the challenges for object tracking: a good object tracker should be able to identify if the object is no longer in the scene and handle its re-appearance.

For comparison purposes we provide results using both evaluation strategies in Table S6; when needed, we re-evaluated using results shared by authors. Our LucidTracker improves over previous results independent of the evaluation strategy.

	Method	# training images	Flow $\mathcal{F}$	DAVIS, mIoU						
				J			F			T
				Mean ↑	Recall ↑	Decay ↓	Mean ↑	Recall ↑	Decay ↓	Mean ↓
	Box oracle [25]	0	✗	45.1	39.7	**-0.7**	21.4	6.7	1.8	**1.0**
	Grabcut oracle [25]	0	✗	67.3	76.9	1.5	65.8	77.2	2.9	34.0
Ignores 1st frame annotation	Saliency	0	✗	32.7	22.6	-0.2	26.9	10.3	**0.9**	32.8
	NLC [13]	0	✓	64.1	73.1	8.6	59.3	65.8	8.6	35.8
	MP-Net [54]	~22.5k	✓	69.7	82.9	5.6	66.3	78.3	6.7	68.6
	Flow saliency	0	✓	70.7	83.2	6.7	69.7	82.9	7.9	48.2
	FusionSeg [23]	~95k	✓	71.5	-	-	-	-	-	-
Uses 1st frame annotation	Mask warping	0	✓	32.1	25.5	31.7	36.3	23.0	32.8	8.4
	FCP [42]	0	✓	63.1	77.8	3.1	54.6	60.4	3.9	28.5
	BVS [33]	0	✗	66.5	76.4	26.0	65.6	77.4	23.6	31.6
	ObjFlow [55]	0	✓	71.1	80.0	22.7	67.9	78.0	24.0	22.1
	STV [60]	0	✓	73.6	-	-	72.0	-	-	-
	VPN [24]	~2.3k	✗	*75.0*	-	-	*72.4*	-	-	*29.5*
	OSVOS [6]	~2.3k	✗	*79.8*	*93.6*	*14.9*	*80.6*	***92.6***	*15.0*	*37.6*
	MaskTrack [25]	~11k	✓	80.3	93.5	8.9	75.8	88.2	9.5	18.3
	LucidTracker	**24~126**	✓	**84.8**	**94.6**	4.3	**82.3**	90.5	7.0	15.8

Table S2: Comparison of segment tracking results on DAVIS benchmark. Numbers in italic are computed based on subsets of DAVIS. Our LucidTracker improves over previous results.

Attribute	Method				
	BVS[33]	ObjFlow [55]	OSVOS [6]	MaskTrack[25]	LucidTracker
Appearance change	0.46	0.54	***0.81***	0.76	0.78
Background clutter	0.63	0.68	*0.83*	0.79	**0.85**
Camera-shake	0.62	0.72	*0.78*	0.78	**0.87**
Deformation	0.7	0.77	*0.79*	0.78	**0.87**
Dynamic background	0.6	0.67	*0.74*	0.76	**0.77**
Edge ambiguity	0.58	0.65	*0.77*	0.74	**0.78**
Fast-motion	0.53	0.55	*0.76*	0.75	**0.80**
Heterogeneous object	0.63	0.66	*0.75*	0.79	**0.83**
Interacting objects	0.63	0.68	*0.75*	0.77	**0.84**
Low resolution	0.59	0.58	***0.77***	**0.77**	0.76
Motion blur	0.58	0.6	*0.74*	0.74	**0.83**
Occlusion	0.68	0.66	*0.77*	0.77	**0.83**
Out-of-view	0.43	0.53	*0.72*	0.71	**0.84**
Scale variation	0.49	0.56	*0.74*	0.73	**0.76**
Shape complexity	0.67	0.69	*0.71*	0.75	**0.81**

Table S3: Attribute evaluation. LucidTracker improves across the bulk of tracking challenges.

H. Additional qualitative results

In this section we provide additional qualitative results for LucidTracker$^-$.

Figure S3 shows qualitative results for the best performing video sequences across three different datasets. Our system handles well camera view changes, object shape deformation, and occlusions.

Failure cases. Figure S4 reports failure cases of LucidTracker$^-$ across three different datasets. For each dataset we show 2 out of 5 worst results (based on average IoU over the video frames).

Video sequence	Method, mIoU				
	BVS[33]	ObjFlow [55]	OSVOS [6]	MaskTrack[25]	LucidTracker
bear	**95.5**	94.6	-	93.1	**95.5**
blackswan	94.3	94.7	94.2	90.3	**95.0**
bmx-bumps	43.4	48.0	-	57.1	**74.2**
bmx-trees	38.2	14.9	55.5	**57.5**	55.0
boat	64.4	**80.8**	-	54.7	74.1
breakdance	50.0	49.6	70.8	76.1	**87.2**
breakdance-flare	72.7	76.5	-	77.6	**93.6**
bus	86.3	68.2	-	89.0	**90.4**
camel	66.9	86.7	85.1	80.1	**94.3**
car-roundabout	85.1	90.0	95.3	**96.0**	**96.0**
car-shadow	57.8	84.6	**93.7**	93.5	90.3
car-turn	84.4	87.6	-	**88.6**	87.8
cows	89.5	91.0	**94.6**	88.2	93.1
dance-jump	74.5	80.4	-	78.8	**83.8**
dance-twirl	49.2	56.7	67.0	84.4	**88.6**
dog	72.3	89.7	90.7	90.8	**95.0**
dog-agility	34.5	86.0	-	78.9	**89.0**
drift-chicane	3.3	17.5	83.5	**86.2**	1.4
drift-straight	40.2	31.4	67.6	56.0	**79.9**
drift-turn	29.9	3.5	-	86.0	**87.1**
elephant	84.9	87.9	-	87.2	**94.2**
flamingo	88.1	87.3	-	79.0	**89.3**
goat	66.1	86.5	88.0	84.5	**88.9**
hike	75.5	93.4	-	93.1	**94.8**
hockey	82.9	85.0	-	83.4	**90.6**
horsejump-high	80.1	86.2	78.0	81.8	**87.1**
horsejump-low	60.1	82.2	-	80.6	**86.7**
kite-surf	42.5	**70.2**	68.6	60.0	64.6
kite-walk	87.0	**85.0**	-	64.5	79.9
libby	77.6	59.4	80.8	77.5	**85.5**
lucia	90.1	89.7	-	**91.1**	90.4
mallard-fly	60.6	55.0	-	57.2	**77.2**
mallard-water	90.7	89.9	-	90.4	**91.4**
motocross-bumps	40.1	48.5	-	59.9	**90.2**
motocross-jump	34.1	59.4	**81.6**	68.3	75.1
motorbike	56.3	47.8	-	56.7	**61.3**
paragliding	87.5	94.7	-	**95.9**	95.6
paragliding-launch	**64.0**	63.7	62.5	62.1	63.7
parkour	75.6	86.1	85.6	88.2	**93.2**
rhino	78.2	89.5	-	91.1	**95.5**
rollerblade	58.8	88.6	-	78.7	**88.8**
scooter-black	33.7	76.5	71.1	82.4	**86.5**
scooter-gray	50.8	29.6	-	82.9	**86.1**
soapbox	78.9	68.9	81.2	89.9	**90.5**
soccerball	84.4	8.0	-	**89.0**	86.8
stroller	76.7	87.7	-	85.4	**90.9**
surf	49.2	**95.6**	-	92.8	94.2
swing	78.4	60.4	-	81.9	**86.9**
tennis	73.7	81.8	-	86.2	**90.7**
train	87.2	91.7	-	90.4	**92.4**
Mean	66.5	71.1	79.8	80.3	**84.8**

Table S4: Per-sequence results on DAVIS.

Video sequence	Method, mIoU				
	BVS[33]	ObjFlow [55]	TRS [63]	MaskTrack[25]	LucidTracker
bird of paradise	89.7	87.1	90.0	**91.5**	88.1
birdfall	65.3	52.9	72.5	68.0	**73.6**
bmx #1	67.1	87.9	86.1	81.2	**91.0**
bmx #2	3.2	4.0	**40.3**	1.0	2.8
cheetah #1	5.4	25.9	61.2	**69.0**	54.8
cheetah #2	9.2	37.2	39.4	18.7	**50.5**
drift #1	68.5	**77.9**	70.7	61.9	66.4
drift #2	32.7	27.4	70.7	**83.6**	80.2
frog	76.1	78.4	80.2	86.2	**90.4**
girl	86.5	84.2	86.4	82.9	**87.4**
hummingbird #1	53.2	67.2	53.0	60.9	**71.9**
hummingbird #2	28.7	68.5	70.5	63.3	**79.2**
monkey	85.7	87.8	83.1	**88.3**	69.2
monkeydog #1	40.5	47.1	74.0	21.6	**79.4**
monkeydog #2	17.1	21.0	39.6	77.1	**81.0**
parachute	93.7	93.3	**95.9**	94.5	93.8
penguin #1	81.6	80.4	53.2	59.6	**94.6**
penguin #2	82.0	83.5	72.9	92.9	**95.2**
penguin #3	78.5	83.9	74.4	89.8	**93.4**
penguin #4	76.4	86.2	57.2	77.2	**93.1**
penguin #5	47.8	**82.3**	63.5	63.9	78.8
penguin #6	84.3	87.3	65.7	88.0	**88.3**
soldier	55.3	**86.8**	76.3	83.2	75.0
worm	65.4	**83.2**	82.4	**83.2**	83.0
Mean	58.4	67.5	69.1	70.3	**77.6**

Table S5: Per-sequence results on SegTrack$_{v2}$.

Method	# training images	1st frame supervision	Flow $\mathcal{F}$	YoutbObjs, mIoU	
				all frames	only frames with object present
FusionSeg [23]	~95k	✗	✓	67.9	71.0
BVS [33]	0	✓	✗	59.7	68.0
N15 [34]	0	✓	✓	-	74.1
ObjFlow [55]	0	✓	✓	70.1	77.6
OSVOS [6]	~2.3k	✓	✗	72.5	78.3
MaskTrack [25]	~11k	✓	✓	72.6	77.8
LucidTracker	**24~126**	✓	✓	**76.2**	**79.8**

Table S6: Comparison of segment tracking results on YouTubeObjects with different evaluation strategies. Our LucidTracker improves over previous results independent of the evaluation strategy.

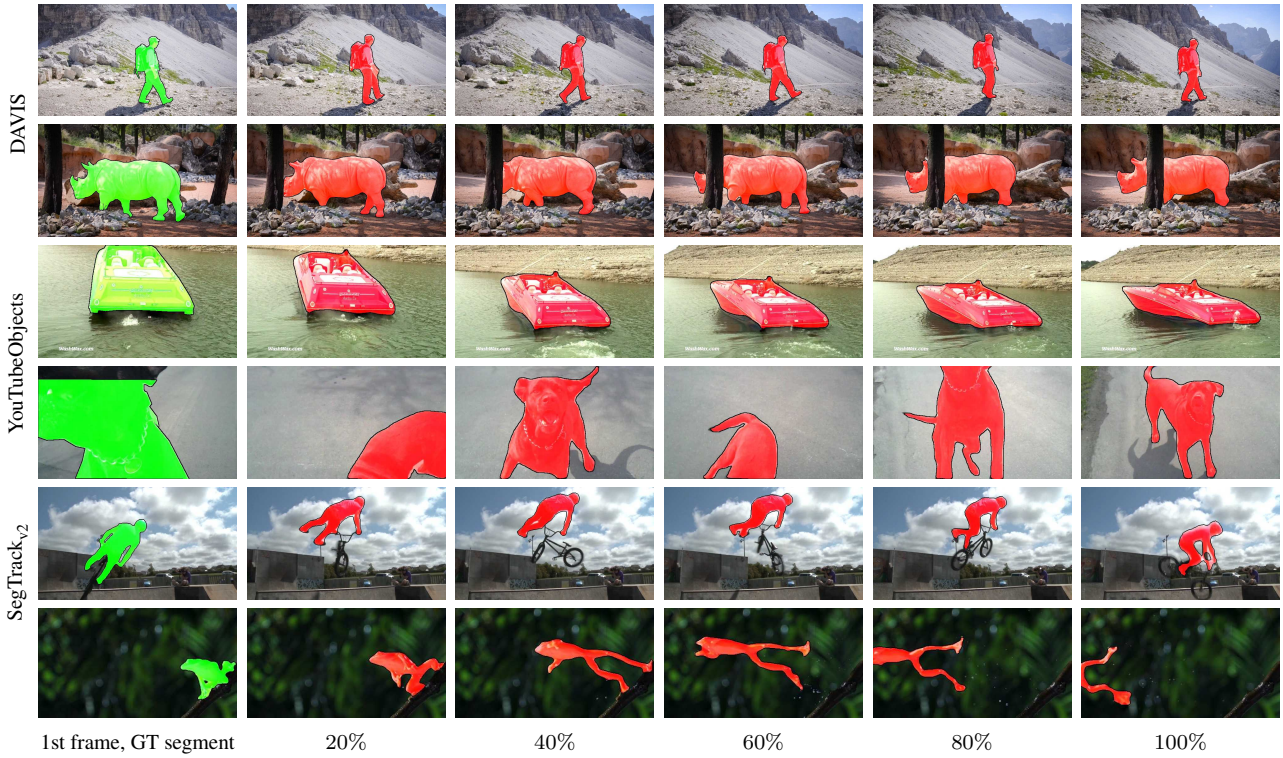

Figure S3: LucidTracker^{-} qualitative results. Frames sampled along the video duration (e.g. 50%: video middle point). For each dataset we show 2 out of the best performing videos (based on mIoU over the video). Our model is robust to various challenges, such as view changes, shape deformations, and occlusions.

Figure S4: LucidTracker^{-} failure cases. Frames sampled along the video duration (e.g. 50%: video middle point). For each dataset we show 2 out of 5 worst results (based on mIoU over the video).